arXiv:1401.6184v2 [cond-mat.stat-mech] 29 Apr 2014

Molecular Dynamics in the Multicanonical Ensemble: Equivalence of Wang–Landau Sampling, Statistical Temperature Molecular Dynamics, and Metadynamics

Christoph Junghans, Danny Perez, and Thomas Vogel*
Theoretical Division T-1, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
(Dated: April 15, 2014)

We show direct formal relationship between the Wang–Landau iteration [PRL 86, 2050 (2001)], metadynamics [PNAS 99, 12562 (2002)] and statistical temperature molecular dynamics (STMD) [PRL 97, 050601 (2006)], the major work-horses for sampling from generalized ensembles. We demonstrate that STMD, itself derived from the Wang–Landau method, can be made indistinguishable from metadynamics. We also show that Gaussian kernels significantly improve the performance of STMD, highlighting the practical benefits of this improved formal understanding.

I. INTRODUCTION

Generalized ensemble methods have become the standard techniques to explore the energy landscape of complex systems [1]. From such samplings, the free energy can be obtained, which provides various thermodynamic insights. The idea of performing Monte Carlo (MC) simulations in non-canonical or extended ensembles goes back a long time. Early milestones include works by Torrie and Valleau [2], who introduced the so-called Umbrella Sampling, Challa and Hetherington [3] who proposed a Gaussian ensemble to interpolate between microcanonical and canonical views of phase transitions in finite systems, and Lyubartsev et al. [4] who simulated an expanded ensemble covering a wide temperature range. Monte Carlo (MC) simulations in the multicanonial (muca) ensemble, first proposed by Berg and Neuhaus [5], exploit the umbrella sampling idea by generating an umbrella in a way that a random walk in energy space is obtained. Later, Hansmann *et al.* [6] extended multicanonial MC to molecular dynamics (MD). Of course, the main technical challenge is the determination of good umbrellas (multicanonial weights) in order to achieve a diffusive behavior in energy space. In a related effort, Wang and Landau (WL) proposed a random-walk algorithm [7] for MC applications, in which the density of states, suitable to calculate multicanonical weights, is estimated on the fly; in fact, over the last decade the WL method has become the most popular tool for this purpose in the MC community [8]. Shortly after, Laio and Parrinello [9] proposed an MD-based method — metadynamics — to fill up basins of the free-energy surface and enhance the exploration of configuration space. Using a different approach, and independently from metadynamics, Kim *et al.* [10] later combined ideas from Hansmann's multicanonial MD with the WL MC algorithm and put forward a method known as statistical temperature molecular dynamics (STMD). Other combinations of WL and MD using the weighted histogram method [11] and/or smoothing of the estimated density of states [12] have been proposed as well.

In this paper we investigate the relationship between WL, STMD and metadynamics. While all these methods are well established in their communities, their precise formal relationships has, to the best of our knowledge, not been thoroughly analysed and, consequently, their development largely proceeds in parallel. We provide a unified formal view of these three methods and give the conditions under which they are equivalent. In particular, we show that STMD and metadynamics produce, on a timestep per timestep basis, *identical* dynamics when using consistent initialization and update schemes. This unified view allows for the transfer of innovation between the different methods and avoids duplication of efforts in different communities.

II. MOLECULAR DYNAMICS IN THE MULTICANONICAL ENSEMBLE

In the muca ensemble, one aims at sampling from a *flat* potential energy distribution $P_{\mathrm{muca}}(U)$, i.e., one requires

$$P_{\mathrm{muca}}(U) \propto g(U)\, w_{\mathrm{muca}}(U) = \mathrm{const}\,., \tag{1}$$

where $g(U)$ is the density of states and $w_{\mathrm{muca}}(U)$ are the multicanonical weights, independently of T. Obviously, for this to be realized, the weights have to take the form

$$w_{\mathrm{muca}}(U) \propto 1/g(U) = \mathrm{e}^{-\ln g(U)} = \mathrm{e}^{-k_{\mathrm{B}}^{-1} S(U)}\,, \tag{2}$$

where $S(U) = k_{\mathrm{B}} \ln g(U)$ is the microcanonical entropy and k_{B} the Boltzmann constant. In the traditional formulation where only configurational degrees of freedom are taken into account, the muca weights can be seen as canonical weights at a temperature T_0 for an effective potential

$$V_{\mathrm{eff}}(U) = T_0\, S(U)\,. \tag{3}$$

* E-mail: tvogel@lanl.gov

The interatomic forces for muca MD simulations are obtained from the gradient of $V_{\rm eff}(U)$:

$$f_i^{\rm muca} = -\frac{{\rm d}V_{\rm eff}(S(U(q_1,\ldots,q_{3n})))}{{\rm d}q_i} = -T_0\,\frac{\partial S}{\partial U}\,\frac{{\rm d}U(q_1,\ldots,q_{3n})}{{\rm d}q_i}\,. \quad (4)$$

Using the definition of the microcanonical temperature:

$$T(U)^{-1} = \partial S(U)/\partial U\,, \quad (5)$$

the multicanonical forces become:

$$f_i^{\rm muca} = \frac{T_0}{T(U)}\,f_i\,, \quad (6)$$

i.e., muca forces differ from the conventional forces, f_i, only by an energy dependent rescaling factor $\propto 1/T(U)$.

Since the multicanonical weights are related to the density of states (Eq. 1), results of a single multicanonical simulation can be reweighted to obtain canonical averages at any temperature. The key difficulty in flat-histogram simulations, on the other hand, is to determine the simulation weights (i.e., the density of states), and many different approaches have been proposed to address that issue, WL being one of the most popular. In WL [7], the density of states $g(U)$ is approximated using a discrete histogram. At each step, the value of the bin of the instantaneous estimator $g'(U,t)$ [13] containing the current energy is updated using a modification factor $f_{\rm WL}$ via

$$\ln\,g'(U_{\rm act},t+\Delta t) = \ln\,g'(U_{\rm act},t) + \ln f_{\rm WL}\,, \quad (7)$$

where 'act' is the actual bin index and t is the MC (or later, MD) simulation time. Conventionally, $\ln f_{\rm WL}$ is initially set to 1 and $\ln\,g'(U,t=0)=0$. Simultaneously, a histogram $H(U)$ of the energy bins visited during the simulation is accumulated. Once $H(U)$ is deemed flat enough, $f_{\rm WL}$ is decreased, e.g., as $f_{\rm WL}\to\sqrt{f_{\rm WL}}$. In this paper, we are mainly concerned with the first iteration, where the dynamics are still strongly biased, but it can be shown that, as $f_{\rm WL}$ tends to 1, the WL method converges to a correct multicanonical sampling [7, 14].

Direct applications of the WL strategy to MD have been attempted [15], however, such efforts were hampered by numerical stability issues introduced by finite-difference differentiation of noisy histograms, requiring the introduction of rather elaborate smoothing procedures [12]. To avoid such complications, Kim *et al.* [10] proposed to directly estimate $T(U)$ (cf. Eqs. 5 and 6) and update $T'(U,t)$, which they refer to as *statistical temperature*, as the MD simulation proceeds *and* to begin from an initially constant temperature $T'(U,t=0)=T_0>0$ instead of a constant entropy as done in WL. This approach allows for a restriction of the sampled temperature range, for example to positive values. Except for that key difference, the STMD scheme is a direct translation of the WL ideas, making muca MD simulations according to Eq. (6) feasible. Applying a central difference approximation to the derivative in Eq. (5), the WL update (Eq. 7) then translates into the following temperature update ($T'(U,t)$ is also a binned, discrete function) in the energy bins next to the currently occupied one:

$$T'(U_{\rm act\pm1},t+\Delta t) = \frac{T'(U_{\rm act\pm1},t)}{1\mp\delta_\beta\,T'(U_{\rm act\pm1},t)}\,, \quad (8)$$

with $\delta_\beta = k_{\rm B}\ln f_{\rm WL}/2\Delta U$ and ΔU being the energy bin width. See Ref. [10] for all details.

Various extensions of this single-bin based update scheme are possible. For example, one can choose any scalable kernel function $\gamma\,k(x/\hat\delta)$ to evolve the entropy estimator $S'(U,t)\propto\ln g'(U,t)$. The update (which can now affect an arbitrarily large energy range) then reads:

$$\ln\,g'(U,t+\Delta t) = \ln\,g'(U,t) + \gamma\,k\left[(U-U_{\rm act})/\hat\delta\right]. \quad (9)$$

This scheme has proven to be particularly useful for Wang–Landau sampling of joint densities of states, i.e., when performing random walks in more than one dimension [16]. The above expression can be cast in terms of an entropy estimator as:

$$S'(U,t) = \gamma\sum_{t^*\le t} k\left[(U-U(t^*))/\hat\delta\right] + S'(U,t=0)\,, \quad (10)$$

where we use the times t^* to index the entropy-update events. Following STMD, assume the initial guess $S'(U,t=0)$ is such that

$$\frac{1}{T'(U,t=0)} = \frac{\partial S'(U,t=0)}{\partial U} = \frac{1}{T_0}\,. \quad (11)$$

Recalling Eq. (4), we then get for the muca forces:

$$\begin{aligned} f_i^{\rm muca\prime}(U,t) &= T_0\,\frac{\partial S'(U,t)}{\partial U}\,f_i \\ &= T_0\left(\frac{\partial}{\partial U}\,\gamma\sum_{t^*\le t}k\left[\frac{U-U(t^*)}{\hat\delta}\right] + \frac{\partial S'(U,t=0)}{\partial U}\right)f_i \\ &= \left(1+\gamma\,T_0\frac{\partial}{\partial U}\sum_{t^*\le t}k\left[\frac{U-U(t^*)}{\hat\delta}\right]\right)f_i\,. \end{aligned} \quad (12)$$

Taking a step back, we can use this last equation to factorize $V_{\rm eff}(U)$ (cf. Eq. 3) into a *sum* of the original potential U and a bias potential V_G: $V_{\rm eff}(U) = U + V_G$. By inspection (cf. Eqs. 4 and 12), we directly get:

$$V_G'(U,t) = \gamma\,T_0\sum_{t^*\le t}k\left[(U-U(t^*))/\hat\delta\right]\,, \quad (13)$$

i.e., with the proper initial conditions, WL/STMD updates are equivalent to the construction of an additive bias potential that takes the form of a simple sum of kernel functions. As we will now show, this procedure is

functionally equivalent to a metadynamics [9] approach with the potential energy as a collective variable. In metadynamics, one also aims at overcoming free energy barriers, allowing for a random walk in the collective-variable space [17]. In order for the system to freely diffuse with respect to the potential energy, the average "metadynamics force" ϕ_F on the collective variable must vanish, i.e., the *free energy* landscape $F_{T_0}(U) = U - T_0 S(U)$ must become flat. To that effect, an additive potential $V_{\rm G}(U)$ is introduced such that

$$\phi_F(U) = \frac{\partial[F_{T_0}(U) + V_{\rm G}(U)]}{\partial U} = 0\,. \qquad (14)$$

Clearly, $V_{\rm G}(U) = -F(U)$ solves Eq. (14), which implies $U + V_{\rm G}(U) = T_0 S(U)$, up to an arbitrary additive constant. Therefore, an energy-based metadynamics simulation simply reduces to a multicanonical MD simulation in U (cf. Eq. 3). In practice, metadynamics starts with the initial guess $V'_{\rm G}(U, t=0) = 0$ for the modifying potential (i.e., also starting the simulation in the canonical ensemble at temperature T_0), which is then gradually updated following a scheme introduced earlier in the energy landscape paving method [18]. Typically, Gaussian kernel functions $k(x/\hat{\delta}) \propto \exp\left[-\frac{1}{2}(x/\hat{\delta})^2\right]$ are used and $V'_{\rm G}(U,t)$ reads:

$$V'_{\rm G}(U,t) = w \sum_{t' \le t} \exp\left[-\frac{(U - U(t'))^2}{2\,\delta U^2}\right], \qquad (15)$$

where w is a tunable constant. The modified interatomic forces are obtained from the gradient of the modified potential $U(q_1, \ldots, q_{3n}) + V_{\rm G}[U(q_1, \ldots, q_{3n}), t]$ [19]:

$$\begin{aligned} f_i^{\rm mod\prime}(U,t) &= \frac{\partial U}{\partial q_i} + \frac{\partial V'_{\rm G}}{\partial U}\frac{\partial U}{\partial q_i} \\ &= f_i \left(1 + \frac{\partial}{\partial U}\, w \sum_{t' \le t} \exp\left[-\frac{(U - U(t'))^2}{2\,\delta U^2}\right]\right), \end{aligned} \qquad (16)$$

which is indeed identical to Eq. (12) when $\gamma\, k(x/\hat{\delta})$ is a Gaussian kernel function with $w = \gamma\, T_0$ and when using the same time sampling points t' and $t^\star$, respectively.

III. RESULTS AND DISCUSSION

During the last decade, there have been multiple independent algorithmic advances in the MC and MD communities that lead to significant improvement of the major generalized ensemble methods, see Refs. [20, 21] for some examples, and the introduction of STMD [10] was a major step in bridging the gap between MC and MD. Our demonstration that STMD and metadynamics can be made identical should further facilitate technological transfers between both communities.

The use of Gaussian kernels, as done in metadynamics, in STMD is the most obvious example of such a transfer. For illustrative purposes (see the Appendix for another example), potential gains are demonstrated using a system consisting of 500 silver atoms at constant particle density $\rho = 0.0585\,\text{Å}^{-3}$, interacting via an embedded-atom potential [22]. We use the stochastic Velocity-Verlet algorithm [23] with a time step of 2 fs and a Langevin thermostat at $T_0 = 3500\,\text{K}$ and apply periodic boundary conditions. We use the original STMD method, where the statistical temperature is updated according to a single-bin update of the entropy (via Eq. 8), and compare with the Gaussian-kernel version where we directly solve Eq. (5). Applying Eq. (9), this leads to the following temperature update:

$$\begin{aligned} T^{-1\prime}(U, t+\Delta t) &= \frac{\partial S'(U, t+\Delta t)}{\partial U} \\ &= k_{\rm B}\,\frac{\partial}{\partial U}\left[\ln g'(U,t) + \gamma\, \mathrm{e}^{-(U-U(t))^2/2\hat{\delta}^2}\right] \\ &= T^{-1\prime}(U,t) - 2\gamma k_{\rm B}\left[(U - U(t))/2\hat{\delta}^2\right] \mathrm{e}^{-(U-U(t))^2/2\hat{\delta}^2}\,. \end{aligned} \qquad (17)$$

We apply Gaussian kernels of different widths, which we measure in units of the energy bin width ΔU used in the original-STMD run. That is, for $\hat{\delta} = n\Delta U/\sqrt{2}$ the kernel function drops to γ/e at the centers of the nth nearest neighbor energy bins. γ takes the role of $\ln f_{\rm WL}$ (cp. Eq. 7) and can be chosen much smaller than for WL simulations [10], we initially set $\gamma = 3.5 \times 10^{-3}$. We furthermore use a cutoff of $10\Delta U$ at both sides of the Gaussian in all cases, but verified that the actual choice of the cutoff does not systematically affect the results (see Appendix for a more detailed discussion and data). The energy-histogram bin width is identical in all cases and the energy histogram itself is always updated by increasing single bins, i.e., the Gaussian kernels are not applied for recording the histogram of visited energies. Also, the flatness criterion is identical for all runs. In Table I we show the average times needed for different runs to fulfill the histogram flatness criterion, i.e., to finish the first WL iteration and, in particular, to visit all energy levels. The result clearly shows that the width of the Gaussian kernel influences how fast the system is driven through energy space, and that wider kernels provide a significant speedup. In Fig. 1 (a), we show time series for the first iteration from two runs, applying the original STMD and a Gaussian kernel run with width $3\Delta U/\sqrt{2}$, respectively. For the latter case, the system moves from

TABLE I. Average times to complete first iteration, i.e., create a flat histogram in a given energy range. Statistical errors were estimated through multiple independent runs.

Method	time in ns
original STMD ($\Delta U = 2\,\text{eV}$)	81.3 ± 27.6
Gaussian kernel, $\hat{\delta} = \Delta U/\sqrt{2}$	88.7 ± 36.6
Gaussian kernel, $\hat{\delta} = 2\,\Delta U/\sqrt{2}$	38.8 ± 18.7
Gaussian kernel, $\hat{\delta} = 3\,\Delta U/\sqrt{2}$	25.9 ± 23.3

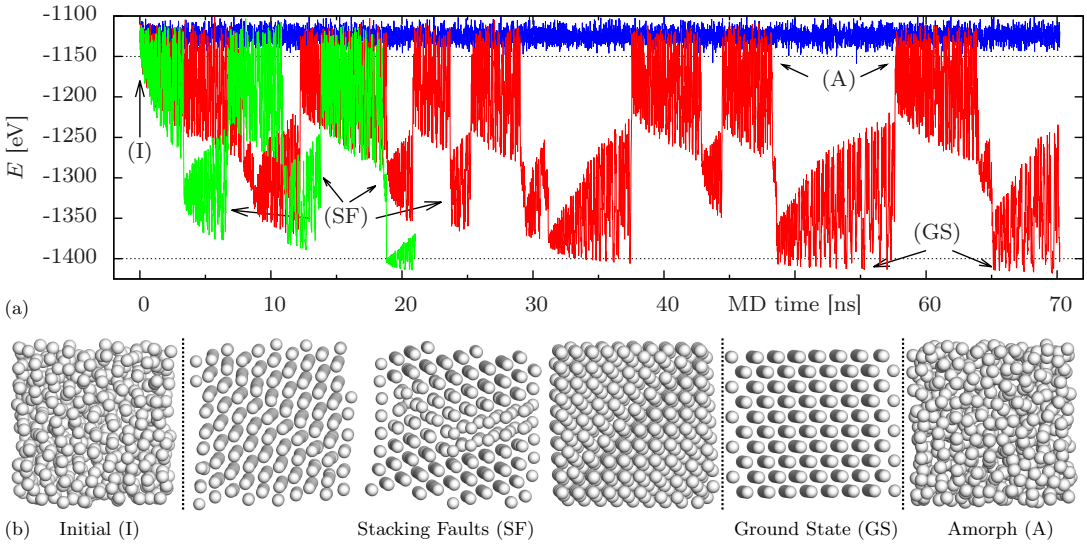


FIG. 1. (a) Time series of the first iteration of original (red) and Gaussian-kernel STMD runs (green), compared to canonical MD at $T = T_0$ (blue). (b) Snapshots of actual atomic configurations sampled during the runs.

the initial (I), amorphous configuration via low-energy crystalline states exhibiting stacking faults (SF) to the perfectly ordered ground state (GS; see Fig. 1 b for visualizations) in just about 20 ns. Concomitantly, extensive thermodynamic information is gathered. Also note that the use of continuous kernel functions, rather than updates of binned estimators, allows, in principle, for an arbitrarily fine-meshed estimation of $T(U)$ without systematically influencing the algorithmic runtime.

Many other improvements can be considered and parallel efforts in the different communities are a common occurrence. For example, it has been shown that the WL energy probability distribution is attracted to the vicinity of the uniform distribution, i.e., that the algorithm converges to the right solution [14]. By introducing a height-reduction scheme for the Gaussian kernels [24] similar statements should be available for MD methods. A more recent development in the metadynamics community concerns adaptive Gaussians [25], where the form of the update to the bias potential depends on local properties of the underlying free-energy surface. Similar ideas of applying different entropy updates in Wang–Landau simulations have circulated [26] and an ad-hoc method for nonuniform binning of energy levels has been recently and independently implemented [27]. To mention a final example, in efforts to develop massively parallel implementations, multiple parallel walkers have been *simultaneously* deployed to update a bias potential in metadynamics [28]. However, systematic errors, unnoticed in Ref. [28], were detected when exactly the same approach was independently applied in the MC community [26]. Joining insights from both studies might lead to further improvements. In fact, a generic parallel scheme based on replica exchanges, which avoids such artificial bias, was recently introduced and applied in both communities [29–32].

IV. SUMMARY

We aim at consolidating the developments in the different areas of generalized ensemble MC and MD sampling by demonstrating that three popular methods, namely Wang–Landau, Statistical Temperature Molecular Dynamics, and Metadynamics, are formally equivalent upon a consistent choice of initial conditions and update rules. Specifically, we have shown that STMD, a translation of the Wang–Landau method into the MD language, augmented by the introduction of kernel updates of the statistical temperature becomes completely identical to metadynamics, i.e., both methods give identical dynamics on a timestep by timestep basis. The focus of this paper is on this explicit equivalence; discussions concerning the overall convergence behavior and analogies in that regard between different strategies in Wang–Landau sampling and metadynamics can be found in the literature, see Refs. [14, 33–35]. We believe that a consistent view of flat-histogram methods as presented here is beneficial to foster transfer of ideas between the respective communities.

ACKNOWLEDGMENTS

We thank A.F. Voter for discussions during the entire project and Y.W. Li, L. Vernon, J. Kim, J.E. Straub, and T. Keyes for critical reading of the manuscript. TV and DP acknowledge funding by Los Alamos National Laboratory's (LANL) Laboratory Directed Research and Development ER program, and CJ by a LANL Director's fellowship. Assigned: LA-UR 13-29519. LANL is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. DOE under Contract DE-AC52-06NA25396.

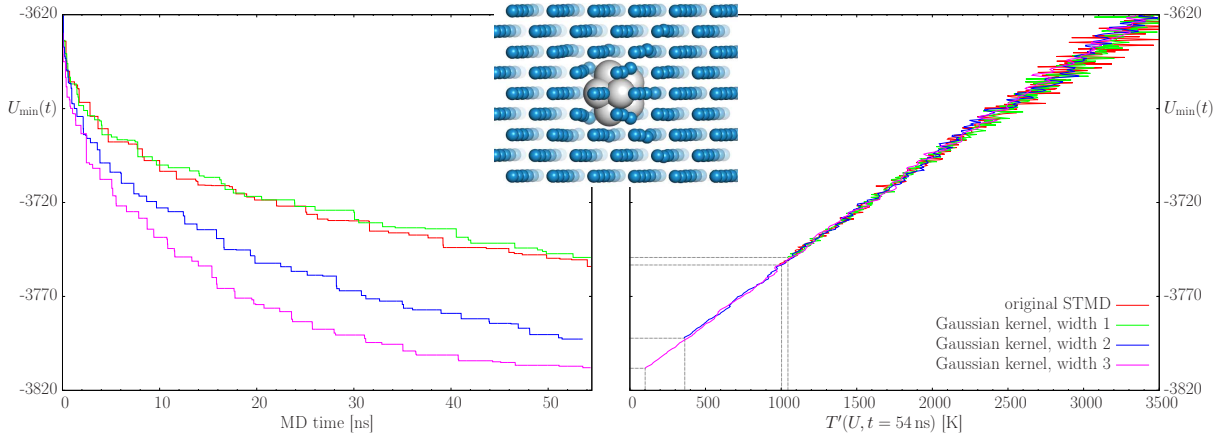


FIG. 2. (Left) Minimal energy values expored by STMD as a function of simulation time for different kernel functions. (Right) The corresponding statistical temperatures after $t \approx 54\,\text{ns}$. Dashed lines mark minimal temperatures corresponding to the minimal energies explored at that time. (Inset) illustration of the physical system in its ground state.

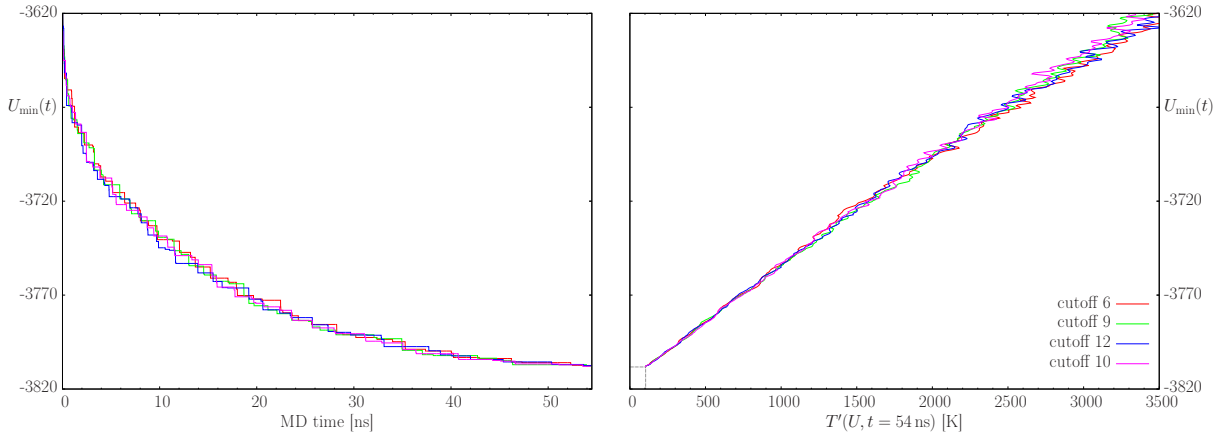


FIG. 3. (Left) Minimal energy values expored by STMD depending on simulation time for different cutoffs of Gaussian kernel functions. (Right) The corresponding statistical temperatures after $t \approx 54\,\text{ns}$.

APPENDIX

Supporting our statements made above, we present here additional data showing the effect of the width of Gaussian kernel functions on the performance of STMD and how a cutoff of those kernels, typically used in practice, affects the results.

Further examples of the speedup due to Gaussian kernel functions

We show in Fig. 2 performance results of STMD simulations on a system composed of a tungsten bcc crystal containing six helium atoms. The clustering of He in W presents an important technological challenge because it can lead to severe microstructural modifications in expected operating conditions of future magnetic-confinement fusion reactors. Indeed, as interstitial He cluster grow, they reach a point where they are able to eject W atoms from the lattice and condense into the resulting vacancies, creating the nuclei of a bubble that can then grow and disrupt the structure of the material (see Refs. [36, 37] for further details and results on that system). The reference temperature is set to $T_0 = 3500\,\text{K}$ (upper temperature boundary of right plot in Fig. 2), which corresponds to a canonical mean energie of $\langle U \rangle_{\text{canonical}} \approx -3620\,\text{eV}$ (upper energy boundary in plots in Fig. 2). When applying wider Gaussian kernel functions (within reasonable limits), the energy and, hence, statistical temperature range is explored much faster, which confirms our claim in the main article. While the original STMD method explores temperatures down to $T \approx 1000\,\text{K}$, temperatures down to $T \approx 100\,\text{K}$ are visited when applying Gaussian kernels of width $\hat{\delta} = 3\Delta U/\sqrt{2}$.

Influence of the cutoff of Gaussian kernel functions

One can reasonably expect that the actual choice of the cutoff of the Gaussian kernel should not affect physical results. Indeed, the original STMD can be seen as the limit

of a short cutoff, and, as shown in Fig. 2 (right panel), its statistical temperature agrees with that of longer cutoff kernels. This is further demonstrated in Fig. 3, where the results presented above were reproduced with different values of the cutoff for Gaussian kernel functions of width $\hat{\delta} = 3\Delta U/\sqrt{2}$ (lowest curve in Fig. 2). Results coincide (right panel) and we observe no notable, systematic difference in the performance in terms of the explored energy and temperature ranges at all times during the runs (left panel). Lowest temperatures explored after $t \approx 54\,\mathrm{ns}$ are all in the range $104 \leq T \leq 107\,\mathrm{K}$. One should however note that using a long cutoff could affect the very low temperatures, as the statistical temperature can approach zero with a finite slope. This discontinuity would be somewhat smoothed out by Gaussian kernels with a large cutoff.

[1] For an overview, see the list of articles citing Refs. [5, 7, 9]. See also, for example, Refs. [20, 21] and references therein.

[2] G. M. Torrie and J. P. Valleau, J. Comp. Phys. **23**, 187 (1977).

[3] M. S. S. Challa and J. H. Hetherington, Phys. Rev. Lett. **60**, 77 (1988); Phys. Rev. A **38**, 6324 (1988).

[4] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).

[5] B. A. Berg and T. Neuhaus, Phys. Lett. B **267**, 249 (1991); Phys. Rev. Lett. **68**, 9 (1992).

[6] U. H. E. Hansmann, Y. Okamoto, and F. Eisenmenger, Chem. Phys. Lett. **259**, 321 (1996).

[7] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001); Phys. Rev. E **64**, 056101 (2001).

[8] In a strict sense, the multicanonical ensemble is an idealized ensemble. Methods creating extended ensembles which eventually converge towards the muca ensemble, such as the original muca recursion, WL sampling, or variations of these, are often referred to as *flat histogram* methods. At finite times, these ensembles are formally different, however, these differences are not critical for our discussion and we will therefore use the terms 'multicanonical' and 'flat-histogram' interchangeably in the following.

[9] A. Laio and M. Parrinello, Proc. Natl. Acad. Sci. U.S.A. **99**, 12562 (2002).

[10] J. Kim, J. E. Straub, and T. Keyes, Phys. Rev. Lett. **97**, 050601 (2006); J. Chem. Phys. **126**, 135101 (2007).

[11] T. Nagasima, A. R. Kinjo, T. Mitsui, and K. Nishikawa, Phys. Rev. E **75**, 066706 (2007).

[12] H. Shimoyama, H. Nakamura, and Y. Yonezawa, J. Chem. Phys. **134**, 024109 (2011).

[13] Primed quantities $y'(x,t)$ will generally refer to instantaneous estimators of $y(x)$ in the following.

[14] C. Zhou and R. Bhatt, Phys. Rev. E **72**, 025701 (2005).

[15] N. Nakajima, H. Nakamura, and A. Kidera, J. Phys. Chem. B **101**, 817 (1997).

[16] C. Zhou, T. C. Schulthess, S. Torbrügge, and D. P. Landau, Phys. Rev. Lett. **96**, 120201 (2006).

[17] U is typically not used as the collective variable in metadynamics; mainly in cases where one aims at estimating the density of states in U [33]. Note that multicanonical and other flat-histogram MC methods have also been widely used with other collective variables as well, with some examples dating even before the introduction of metadynamics. One could mention the bond parameter of Potts-like models [38], the Parisi overlap parameter for spin glasses [39], or interaction parameters in a polymer model [40] as examples. In general, our discussion is independent of the actual choice of this variable, but we use U for clarity.

[18] U. H. E. Hansmann and L. T. Wille, Phys. Rev. Lett. **88**, 068105 (2002).

[19] A. Laio and F. L. Gervasio, Rep. Prog. Phys. **71**, 126601 (2008).

[20] A. Mitsutake, Y. Sugita, and Y. Okamoto, Pept. Sci. **60**, 96 (2001).

[21] S. Singh, M. Chopra, and J. J. de Pablo, Ann. Rev. Chem. Biomol. Eng. **3**, 369 (2012).

[22] P. L. Williams, Y. Mishin, and J. C. Hamilton, Model. Simul. Mater. Sci. Eng. **14**, 817 (2006).

[23] S. Melchionna, J. Chem. Phys. **127**, 044108 (2007).

[24] A. Barducci, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **100**, 020603 (2008).

[25] D. Branduardi, G. Bussi, and M. Parrinello, J. Chem. Theory Comput. **8**, 2247 (2012).

[26] J. Yin and D. Landau, Comput. Phys. Commun. **183**, 1568 (2012).

[27] Y. W. Koh, H. K. Lee, and Y. Okabe, Phys. Rev. E **88**, 053302 (2013).

[28] P. Raiteri, A. Laio, F. L. Gervasio, C. Micheletti, and M. Parrinello, J. Phys. Chem. B **110**, 3533 (2006).

[29] J. Kim, T. Keyes, and J. E. Straub, J. Chem. Phys. **130**, 124112 (2009).

[30] J. Kim, J. E. Straub, and T. Keyes, J. Phys. Chem. B **116**, 8646 (2012).

[31] T. Vogel, Y. W. Li, T. Wüst, and D. P. Landau, Phys. Rev. Lett. **110**, 210603 (2013).

[32] L. Gai, T. Vogel, K. A. Maerzke, C. R. Iacovella, D. P. Landau, P. T. Cummings, and C. McCabe, J. Chem. Phys. **139**, 054505 (2013).

[33] C. Micheletti, A. Laio, and M. Parrinello, Phys. Rev. Lett. **92**, 170601 (2004).

[34] R. E. Belardinelli, S. Manzi, and V. D. Pereyra, Phys. Rev. E **78**, 067701 (2008).

[35] Y. Crespo, F. Marinelli, F. Pietrucci, and A. Laio, Phys. Rev. E **81**, 055701 (2010).

[36] D. Perez, T. Vogel, and B. P. Uberuaga, Phys. Rev. B (2014), submitted.

[37] T. Vogel and D. Perez, Phys. Procedia (2014), in press; arXiv:1404.1407.

[38] C. Chatelain, B. Berche, W. Janke, and P.-E. Berche, Nucl. Phys. B **719**, 275 (2005).

[39] B. A. Berg and W. Janke, Phys. Rev. Lett. **80**, 4771 (1998).

[40] J. Luettmer-Strathmann, F. Rampf, W. Paul, and K. Binder, J. Chem. Phys. **128**, 064903 (2008).